arXiv:1608.00990v1 [stat.CO] 2 Aug 2016

Barrett: out-of-core processing of MultiNest output

Sebastian Liem

GRAPPA, University of Amsterdam, Netherlands

`barrett` is a Python package for processing and visualising statistical inferences made using the nested sampling algorithm `MultiNest`. The main differential feature from competitors are full out-of-core processing allowing `barrett` to handle arbitrarily large datasets. This is achieved by using the HDF5 data format.

`MultiNest` [1, 2] is a nested sampling algorithm widely used to statistical inferences in various astrophysical and particle physics contexts. The output produced by `MultiNest` are plain text files which in turn has to be processed in order to produce and visualise the desired marginalised posteriors, profile likelihood maps and sundry statistical results. There are several packages available to process and visualise the `MultiNest` output [3, 4, 5]. They all, however, assumes that the size of the output to be processed is smaller than the available system memory. But, for large dimensional problem, especially when doing a profile likelihood analysis, the size of that output can become $\mathcal{O}(10 - 100$ GB) which violates that assumption. This paper introduces `barrett` which is a Python package that can handle such large datasets.

Marginalisation and profiling are two ways of reducing the dimensionality of the posterior and likelihood respectively. This is done either to eliminate uninteresting, i.e. nuisance, parameters or to reduce the dimensionality to one or two dimension in order to be able to visualise the results. Denoting the parameters of interests as δ, and the ones to be eliminated as ξ the marginalised posterior and profile likelihood can be defined as:

$$p(\delta|\text{data}) = \int p(\delta, \xi|\text{data})\, \mathrm{d}\xi \qquad L(\text{data}|\delta) = \sup_{\xi} L(\text{data}|\delta, \xi)$$

where $p(\delta, \xi|\text{data})$ and $L(\text{data}|\delta, \xi)$ is the full posterior and likelihood respectively. In the context of the `MultiNest` output, the operation underlying both the marginalisation of posteriors and the profiling of likelihoods is binning. The output is a list of points in the scanned parameter space with an associated posterior

weight and likelihood. To profile or marginalise one define bins in the parameters of interest, and group every point into these bins. The marginalised posterior in each bin then is the sum of the posterior weights of all the points in each bin. The profile likelihood in each bin is the largest likelihood of all the points in each bin.

`barrett` make use of the fact that binning is an operation that can be done in parts. The `MultiNest` output is divided into parts small enough to fit system memory. Each part is then read into memory in turn and binned. The partial results can easily be combined by either summing them (marginalisation) or taking the bin maxima (profiling).

Besides the usual components of scientific Python (`numpy`, `scipy`, and `matplotlib` [6, 7, 8]) the only dependency is the `h5py` package [9] which provides the Python interface to the HDF5 [10] data format. `barrett` uses HDF5 files because it offers two important advantages over the plain text file produced by `MultiNest`. First, HDF5 is a binary format which takes much less disk space when compared to plain text. Secondly, the HDF5 data format allows one to 'chunk' the data, i.e. divides it into parts while still storing it in the same file. The library then only loads the chunks accessed into memory. While the first advantage is a nice bonus it is the second that is essential for `barrett` allowing it to perform the out-of-core binning described in the paragraph above.

1 Installation

`barrett` is available on PyPI and can therefor be installed as easily as

```
$ pip install barrett
```

Alternately the code is available at `https://github.com/sliem/barrett` where any issues or bugs can also be reported.

2 Description

`barrett` has four submodules; `data`, `posterior`, `profilelikelihood`, and `util`.

`barrett.data` implements methods for modifying data (e.g. log, change units) or calculate depended variables (e.g. mean squark mass) in an out-of-core fashion.

`barrett.posterior` defines two classes `oneD` and `twoD` for 1D and 2D marginalised posteriors respectively. Each class performs the marginalisation on initialisation, and implements a plotting method to visualise the results. The `twoD` class also implements a method for calculating the credibility regions.

`barrett.profilelikelihood` also defines two classes named `oneD` and `twoD` but for the 1D and 2D profile likelihoods. The profiling is done on instance initialisation, and the classes define plotting methods. The `twoD` class can also calculate the desired confidence regions.

`barrett.util` contains, as the name suggest, functions of varying utility. Perhaps, the most useful one is `convert_chain()` which converts from the plain text output of `MultiNest` to the HDF5 format used by `barrett`.

The default `plot()` methods are functional but, by design, very basic. Rather than putting another layer of API on top of `matplotlib` `barrett` encourage the user to directly work with `matplotlib`. This can be done either by directly accessing the variables in the classes or by subclassing them and override the `plot()` method.

3 Example

In the following two examples `barrett` will be use to marginalise a posterior two dimensions and then visualise it. The details of the physics isn't of interest, but the parameters used are the dark matter mass and relic density. The results of the example codes are in Fig. 1. The first example shows the basic usage of `barrett`.

```
import barrett.posterior as posterior
import matplotlib.pyplot as plt

h5file = 'RD.h5'
xvar = 'log(m_{\chi})'
yvar = 'Omega_{\chi}h^2'
P = posterior.twoD(h5file, xvar, yvar,
                   xlimits=(0,3), ylimits=(0.0, 0.2),
                   xbins=30, ybins=30)

fig, ax = plt.subplots(ncols=1, nrows=1)
P.plot(ax)

ax.set_xlabel('log10(mDM / 1 GeV)')
ax.set_ylabel('oh^2')
fig.savefig('ex1.png', dpi=200, bbox_inches='tight')
```

The second example shows how to customise the plot by subclassing and overriding the `plot()` method.

```
class myTwoD(posterior.twoD):
    def plot(self, xlabel, ylabel, fig_path):
        fig, ax = plt.subplots(ncols=1, nrows=1)

        # plot heatmap of the posterior.
        posterior.twod.plot(self, ax, levels=none)

        # plot one and two sigma contours.
        x, y = np.meshgrid(self.xcenters, self.ycenters)
        levels = np.append(self.credibleregions([0.95, 0.68]),
                           self.pdf.max())
        ax.contour(x, y, self.pdf, levels=levels, colors='k')

        ax.set_xlabel(xlabel)
        ax.set_ylabel(ylabel)
        fig.savefig(fig_path, dpi=200, bbox_inches='tight')

P = myTwoD(h5file, xvar, yvar,
           xlimits=(0,3), ylimits=(0.0, 0.2),
           xbins=30, ybins=30)
P.plot('log10(mDM / 1 GeV)', 'oh^2', 'ex2.png')
```

Additional examples can be found on `https://github.com/sliem/barrett`.

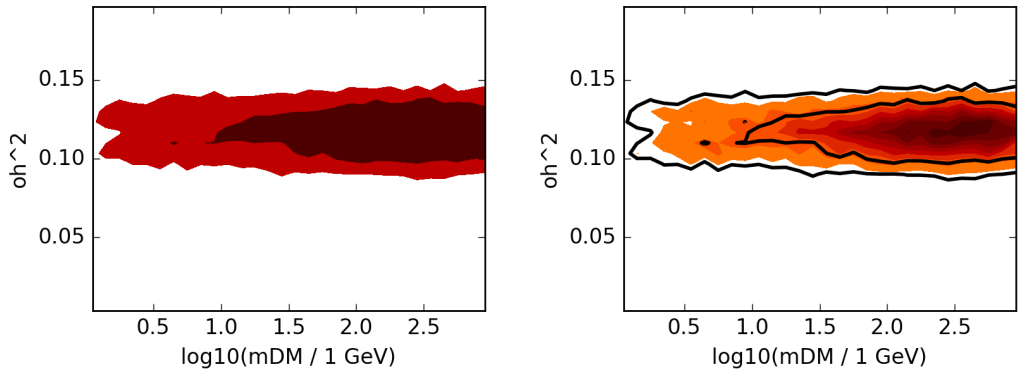

(a) default plot method

(b) custom plot method

Figure 1: The resulting plots of the marginalised posterior using the two example code snippets.

References

[1] F. Feroz and M. P. Hobson, *Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis*, *Mon. Not. Roy. Astron. Soc.* **384** (2008) 449, [arXiv:0704.3704].

[2] F. Feroz, M. P. Hobson, and M. Bridges, *MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics*, *Mon. Not. Roy. Astron. Soc.* **398** (2009) 1601–1614, [arXiv:0809.3437].

[3] A. Fowlie and M. H. Bardsley, *Superplot: a graphical interface for plotting and analysing MultiNest output*, arXiv:1603.00555.

[4] Buchner, J., Georgakakis, A., Nandra, K., Hsu, L., Rangel, C., Brightman, M., Merloni, A., Salvato, M., Donley, J., and Kocevski, D., *X-ray spectral modelling of the agn obscuring region in the cdfs: Bayesian model selection and catalogue*, *A&A* **564** (2014) A125.

[5] A. Lewis, *Getdist*, 2015–. https://github.com/cmbant/getdist/ [Online; accessed 2016-07-31].

[6] S. van der Walt, S. C. Colbert, and G. Varoquaux, *The numpy array: A structure for efficient numerical computation*, *Computing in Science Engineering* **13** (March, 2011) 22–30.

[7] E. Jones, T. Oliphant, P. Peterson, et al., *SciPy: Open source scientific tools for Python*, 2001–. http://www.scipy.org/ [Online; accessed 2016-07-31].

[8] J. D. Hunter, *Matplotlib: A 2d graphics environment*, *Computing in Science Engineering* **9** (May, 2007) 90–95.

[9] A. Collette, *HDF5 for Python*, 2008–. http://www.h5py.org/ [Online; accessed 2016-07-31].

[10] The HDF group, *Hierarchical Data Format, version 5*, 1997–. http://www.hdfgroup.org/HDF5/, [Online; accessed 2016-07-31].